Zeroth-Order Optimization for Composite Problems with Functional Constraints

Zichong Li[1], Pin-Yu Chen[2], Sijia Liu[3], Songtao Lu[2], Yangyang Xu[1]

[1]Department of Mathematical Sciences, Rensselaer Polytechnic Institute
[2]IBM Research
[3]Department of Computer Science and Engineering, Michigan State University
liz19@rpi.edu, Pin-Yu.Chen@ibm.com, liusiji5@msu.edu, songtao@ibm.com, xuy21@rpi.edu

Abstract

In many real-world problems, first-order (FO) derivative evaluations are too expensive or even inaccessible. For solving these problems, zeroth-order (ZO) methods that only need function evaluations are often more efficient than FO methods or sometimes the only options. In this paper, we propose a novel zeroth-order inexact augmented Lagrangian method (ZO-iALM) to solve black-box optimization problems, which involve a composite (i.e., smooth+nonsmooth) objective and functional constraints. This appears to be the first work that develops an iALM-based ZO method for functional constrained optimization and meanwhile achieves query complexity results matching the best-known FO complexity results up to a factor of variable dimension.With an extensive experimental study, we show the effectiveness of our method. The applications of our method span from classical optimization problems to practical machine learning examples such as resource allocation in sensor networks and adversarial example generation.

Introduction

In many practical optimization problems such as black-box attack (Chen et al. 2017), we only have access to zeroth-order (ZO) function values but no access to first-order (FO) or higher order derivatives (Liu et al. 2020a). In this paper, we consider *nonconvex* problems with *possibly nonconvex* constraints:

$$f_0^* := \min_{\mathbf{x}\in\mathbb{R}^d} \left\{ f_0(\mathbf{x}) := g(\mathbf{x}) + h(\mathbf{x}), \text{ s.t. } \mathbf{c}(\mathbf{x}) = \mathbf{0} \right\}, \quad (1)$$

where g is smooth but possibly nonconvex, $\mathbf{c} = (c_1, \ldots, c_l) : \mathbb{R}^d \to \mathbb{R}^l$ is a vector function with continuously differentiable components, and h is closed convex but possibly nonsmooth and has a coordinate structure, i.e., $h(\mathbf{x}) = \sum_{i=1}^d h(x_i)$. This formulation follows from (Li et al. 2021), except that in this paper, only the function evaluations of g and $\mathbf{c}$, but not their gradients, are accessible. Such a formulation includes a large class of nonlinear constrained problems. We remark that an inequality constraint $t(\mathbf{x}) \leq 0$ can be equivalently formulated as an equality constraint $t(\mathbf{x})+s = 0$ by enforcing the nonnegativity of s, with equivalent stationarity conditions (c.f. (Li et al. 2021)). Note in (1), the inclusion of a coordinate-separable constraint set $\mathcal{X}$ is equivalent to the addition of $I_{\mathcal{X}}$ to the nonsmooth term h in the objective f_0, where $I_{\mathcal{X}}(\mathbf{x}) = 0$ if $\mathbf{x} \in \mathcal{X}$ and $+\infty$ otherwise.

Contributions

Our contributions are three-fold. First, we design a zeroth-order accelerated proximal coordinate update (ZO-APCU) method for solving coordinate-structured strongly-convex composite (i.e., smooth+nonsmooth) problems. ZO-APCU appears to be the first PCU method *with acceleration* by just using function values. It can be viewed as a ZO variant of the APCG in (Lin, Lu, and Xiao 2014). To solve black-box optimization in the form of (1), we propose a novel zeroth-order inexact augmented Lagrangian method (ZO-iALM), by using ZO-APCU to design a zeroth-order inexact proximal point method (ZO-iPPM) to approximately solve each ALM subproblem. Though any ZO method can be used as a subroutine in the iALM, the use of ZO-iPPM with the developed ZO-APCU is crucial to yield best-known query complexity results and also good numerical performance, as we demonstrate in the experiments.

Second, query complexity analysis is conducted on the proposed methods. We show that ZO-APCU needs $O(d\sqrt{\kappa}\log\frac{1}{\varepsilon})$ queries to produce an ε-stationary point of a d-dimensional strongly-convex composite problem with a condition number κ. The ZO-iPPM has an $\tilde{O}(d\varepsilon^{-2})$ complexity to give an ε-stationary point of a nonconvex composite problem. On solving (1) that satisfies a regularity condition, the ZO-iALM has an $\tilde{O}(d\varepsilon^{-3})$ overall complexity to produce an ε-KKT point, and the complexity can be reduced to $\tilde{O}(d\varepsilon^{-\frac{5}{2}})$ if the constraints are affine. All our complexity results are (near) optimal or the best known. To the best of our knowledge, complexity of ZO methods on nonconvex functional constrained problem (1) has not been studied in the literature, thus our $\tilde{O}(d\varepsilon^{-3})$ result is completely new.

Third, we use a coordinate gradient estimator while implementing the core solver ZO-APCU. To be able to yield high-accuracy solutions, we give a multi-point coordinate-wise gradient estimator and analyze its error bound. Under the j-th order smoothness assumption for some $j \in \mathbb{Z}^+$, we show that the error of a $\max\{2, 2(j-1)\}$-point coordinate-wise gradient estimator is upper bounded by $O(a^j)$, where a is the sampling radius. This result is meaningful and important to yield high accuracy, because in practice a cannot be too small, otherwise round-off errors will dominate.

Overall, we conduct a comprehensive study on ZO methods on solving nonconvex functional constrained black-box optimization, from multiple perspectives including complexity analysis, gradient estimator, and significantly improved performance on practical machine learning tasks and classical optimization problems.

Related Works

In this subsection, we review previous works on the inexact augmented Lagrangian methods (iALMs) in the usual FO settings and the zeroth-order methods (ZOMs).

The iALM is one of the most common methods for solving constrained problems. It alternatingly updates the primal variable by inexactly minimizing the augmented Lagrangian (AL) function and the Lagrangian multiplier by dual gradient ascent. If the multiplier is fixed to zero, then iALM reduces to a standard penalty method, which usually has a worse practical performance than iALM. Previous works on iALM assume that FO derivatives of the objective function can be evaluated, therefore use FOMs to inexactly minimize the AL function in each primal update.

For convex nonlinear constrained problems, the iALM in (Li and Xu 2021) and the proximal-iALM in (Li and Qu 2021) can produce an ε-KKT point with $O(\varepsilon^{-1}|\log\varepsilon|)$ gradient evaluations, and the AL-based FOMs in (Xu 2021b,a; Ouyang et al. 2015; Li and Qu 2021; Nedelcu, Necoara, and Tran-Dinh 2014) can produce an ε-optimal solution with $O(\varepsilon^{-1})$ complexity. For strongly-convex problems, the complexity results can be respectively reduced to $O(\varepsilon^{-\frac{1}{2}}|\log\varepsilon|)$ for an ε-KKT point and $O(\varepsilon^{-\frac{1}{2}})$ for an ε-optimal solution, e.g., (Li and Xu 2021; Li and Qu 2021; Xu 2021b; Nedelcu, Necoara, and Tran-Dinh 2014; Necoara and Nedelcu 2014).

For nonconvex problems with nonlinear convex constraints, when Slater's condition holds, $\tilde{O}(\varepsilon^{-\frac{5}{2}})$ complexity results have been obtained by the AL or quadratic-penalty based FOMs in (Li and Xu 2021; Lin, Ma, and Xu 2019) and the proximal ALM in (Melo, Monteiro, and Wang 2020a). When a regularity condition (see Assumption 5 below) holds, the ALM in (Li et al. 2021) achieves $\tilde{O}(\varepsilon^{-\frac{5}{2}})$ complexity for nonconvex problems with nonlinear convex constraints and $\tilde{O}(\varepsilon^{-3})$ complexity for problems with nonconvex constraints.

When gradients of the objective function are unavailable, ZOMs are the only tools available. Previous ZO works mainly focus on problems *without* nonlinear functional constraints. Many existing ZOMs are modified from some gradient descent type FOMs, replacing the exact gradient $\nabla f(\mathbf{x})$ by some gradient estimator $\tilde{\nabla} f(\mathbf{x})$. In the next section, we briefly review some existing gradient estimation frameworks including random search and finite difference. A more detailed overview can be found in (Liu et al. 2020a) for ZOMs.

Notations, Definitions, and Assumptions

We use $\|\cdot\|$ for the Euclidean norm of a vector and the spectral norm of a matrix. $[n]$ denotes the set $\{1,\dots,n\}$. We use $\tilde{O}$ to suppress all log terms of ε from the big-O notation. We denote $J_c(\mathbf{x})$ as the Jacobian of $\mathbf{c}$ at $\mathbf{x}$. The distance between a vector $\mathbf{x}$ and a set $\mathcal{X}$ is denoted as $\operatorname{dist}(\mathbf{x},\mathcal{X}) = \min_{\mathbf{y}\in\mathcal{X}}\|\mathbf{x}-\mathbf{y}\|$. For a function f, we use ∂f to denote the subdifferetial of f. For a differentiable function f, we use $\tilde{\nabla} f$ as an estimator of the gradient ∇f. The AL function of (1) is

$$\mathcal{L}_\beta(\mathbf{x},\mathbf{y}) = f_0(\mathbf{x}) + \mathbf{y}^\top \mathbf{c}(\mathbf{x}) + \tfrac{\beta}{2}\|\mathbf{c}(\mathbf{x})\|^2, \tag{2}$$

where $\beta > 0$ is the penalty parameter, and $\mathbf{y}\in\mathbb{R}^l$ is the multiplier or the dual variable.

Definition 1 (ε-KKT point) *A point $\mathbf{x}\in\mathbb{R}^d$ is an ε-KKT point of* (1) *if there is $\mathbf{y}\in\mathbb{R}^l$ such that*

$$\|\mathbf{c}(\mathbf{x})\| \le \varepsilon, \quad \operatorname{dist}\left(\mathbf{0}, \partial f_0(\mathbf{x}) + J_c^\top(\mathbf{x})\,\mathbf{y}\right) \le \varepsilon. \tag{3}$$

Definition 2 (k-smoothness) *For some $k \ge 1$, we say f is M_k k-smooth, if the k-th derivative of f is M_k Lipschitz continuous.*

Remark 1 *Letting $k=1$ above corresponds to the standard smoothness assumption.*

Definition 3 (coordinate k-smooth) *For some $k \ge 1$, we say f is M_k coordinate k-smooth, if the partial function $F_i(\mathbf{x}_i) := f(\mathbf{x}_{<i},\mathbf{x}_i,\mathbf{x}_{>i})$ is M_k k-smooth, $\forall\, i\in[d]$, where $\mathbf{x}_{<i} := (x_1,\dots,x_{i-1})$ and $\mathbf{x}_{>i} := (x_{i+1},\dots,x_d)$ are fixed.*

Remark 2 *If f is M_k k-smooth, then it must be M_k^c coordinate k-smooth with some $M_k^c \le M_k$.*

Definition 4 (ρ-weakly convex) *A function f is ρ-weakly convex if $f + \frac{\rho}{2}\|\cdot\|^2$ is convex.*

Remark 3 *A function that is L-smooth is also L-weakly convex. However, its weak convexity constant can be much smaller than its smoothness constant.*

Throughout this paper, we make the following assumptions.

Assumption 1 (smoothness and weak convexity) *In* (1)*, g is L_0-smooth and ρ_0-weakly convex. For each $j\in[l]$, c_j is L_j-smooth and ρ_j-weakly convex.*

Assumption 2 (bounded domain) *In* (1)*, h is closed convex with a compact domain, i.e.,*

$$D := \max_{\mathbf{x},\mathbf{x}'\in\operatorname{dom}(h)} \|\mathbf{x}-\mathbf{x}'\| < \infty, \tag{4a}$$

$$D_i := \max_{\substack{\mathbf{x},\mathbf{x}'\in\operatorname{dom}(h)\\ \mathbf{x}_{[d]\setminus i}=\mathbf{x}'_{[d]\setminus i}}} \|\mathbf{x}-\mathbf{x}'\| \le D, \forall i\in[d], \tag{4b}$$

where $\mathbf{x}_{[d]\setminus i} = (\mathbf{x}_1,\dots,\mathbf{x}_{i-1},\mathbf{x}_{i+1},\dots,\mathbf{x}_d)$.

Due to the page limit, all proofs in this paper are given in the supplementary material.

Multi-point Gradient Estimator

In this section, we provide backgrounds on gradient estimators and propose the zeroth-order multi-point coordinate gradient estimator.

Backgrounds on Gradient Estimators

Let a denote the *sampling radius* (also called the *smoothing parameter*) of a random gradient estimator, and $\mathbf{u} \sim p$ denote some random direction sampled from a distribution p. Denote f_a as the *smoothed* version of f defined as $f_a(\mathbf{x}) := \mathbb{E}_{\mathbf{u}\sim p'}[f(\mathbf{x} + a\mathbf{u})]$, where p' is a certain distribution determined by p. All random gradients in this subsection are biased with respect to ∇f but unbiased with respect to ∇f_a, satisfying $\mathbb{E}_{\mathbf{u}\sim p}[\tilde{\nabla} f(\mathbf{x})] = \nabla f_a(\mathbf{x})$,

The 1-*point random gradient estimator* of f has the form

$$\tilde{\nabla} f(\mathbf{x}) := \tfrac{\phi(d)}{a} f(\mathbf{x} + a\mathbf{u})\mathbf{u}, \tag{5}$$

where $\phi(d)$ is a dimension-dependent factor given by the distribution of $\mathbf{u}$. If $p = \mathcal{N}(\mathbf{0}, \mathbf{I})$, then $\phi(d) = 1$; if $p = \mathcal{U}(\mathcal{S}(\mathbf{0}, \mathbf{I}))$ is the uniform distribution on the unit sphere, then $\phi(d) = d$. In practice, the 1-point estimator in (5) is not commonly used due to high variance (Flaxman, Kalai, and McMahan 2004). This motivates the 2-*point random gradient estimator* (Nesterov and Spokoiny 2017; Duchi et al. 2015)

$$\tilde{\nabla} f(\mathbf{x}) := \tfrac{\phi(d)}{a}(f(\mathbf{x} + a\mathbf{u}) - f(\mathbf{x}))\mathbf{u}, \tag{6}$$

where $\mathbb{E}_{\mathbf{u}\sim p}[\mathbf{u}] = \mathbf{0}$ is required for unbiasedness to hold. The 2-point estimator has the following upper bound of expected estimation error (Berahas et al. 2021; Liu et al. 2018)

$$\mathbb{E}[\|\tilde{\nabla} f(\mathbf{x}) - \nabla f(\mathbf{x})\|] = O(\sqrt{d})\|\nabla f(\mathbf{x})\| + O\left(\tfrac{ad^{1.5}}{\phi(d)}\right). \tag{7}$$

Note that the $O(\sqrt{d})\|\nabla f(\mathbf{x})\|$ term in (7) does not vanish even if $a \to 0$. Mini-batch sampling can be used to reduce the estimation error, leading to the *multi-point random gradient estimator* (Duchi et al. 2015; Liu et al. 2018)

$$\tilde{\nabla} f(\mathbf{x}) := \frac{\phi(d)}{a}\sum_{i=1}^{b}(f(\mathbf{x} + a\mathbf{u}_i) - f(\mathbf{x}))\mathbf{u}_i, \tag{8}$$

where b is the mini-batch size, and $\{\mathbf{u}_i\}_{i=1}^b$ are random directions drawn from some zero-mean distribution p. The multi-point estimator has the improved error bound (Berahas et al. 2021)

$$\begin{aligned}&\mathbb{E}[\|\tilde{\nabla} f(\mathbf{x}) - \nabla f(\mathbf{x})\|]\\ &= O\left(\sqrt{\frac{d}{b}}\right)\|\nabla f(\mathbf{x})\| + O\left(\frac{ad^{1.5}}{\phi(d)b}\right) + O\left(\frac{ad^{0.5}}{\phi(d)}\right).\end{aligned}$$

To further reduce the estimation error, one can use *coordinate-wise gradient* which requires $O(d)$ queries per gradient estimate. Existing works use *forward difference* $\tilde{\nabla} f(\mathbf{x}) := \frac{1}{a}\sum_{i=1}^d (f(\mathbf{x} + a\mathbf{e}_i) - f(\mathbf{x}))\mathbf{e}_i$ or *central difference* $\tilde{\nabla} f(\mathbf{x}) := \frac{1}{2a}\sum_{i=1}^d (f(\mathbf{x} + a\mathbf{e}_i) - f(\mathbf{x} - a\mathbf{e}_i))\mathbf{e}_i$ as the coordinate gradient, where $\mathbf{e}_i$ is the ith basis vector. Under the standard smoothness assumption, both forward difference and central difference have the error bounds (Kiefer, Wolfowitz et al. 1952; Berahas et al. 2021; Lian et al. 2016)

$$\begin{aligned}\mathbb{E}[|\tilde{\nabla}_i f(\mathbf{x}) - \nabla_i f(\mathbf{x})|] &= O(a),\\ \mathbb{E}[\|\tilde{\nabla} f(\mathbf{x}) - \nabla f(\mathbf{x})\|] &= O\left(a\sqrt{d}\right).\end{aligned}$$

Zeroth-order Multi-point Coordinate Gradient Estimator

In this subsection, assuming f to be coordinate p-smooth, we construct the zeroth-order multi-point coordinate gradient estimator (ZO-MCGE) $\tilde{\nabla}_i f(\mathbf{x})$ with $p = \max\{2(j-1), 2\}$ function value queries at $\mathbf{x} + \frac{p}{2}a\mathbf{e}_i, \dots, \mathbf{x} + a\mathbf{e}_i, \mathbf{x} - a\mathbf{e}_i, \dots, \mathbf{x} - \frac{p}{2}a\mathbf{e}_i$, and analyze its error bound. The main difference between our proposed ZO-MCGE and the estimators in the previous subsection is that the use of multi-point function evaluation allows for a better control for the gradient estimation error. We observe numerically that using more points in the gradient estimator enables us to reach a higher accuracy; see the logistic regression experiment in the Appendix.

The following lemma directly follows from the coordinate j-smoothness of f.

Lemma 1 *Assume f is M_j coordinate j-smooth. Let $\nabla_i^l f(\mathbf{x}) := \frac{\partial^l f(\mathbf{x})}{(\partial x_i)^l}$ be the l-th order derivative of f at $\mathbf{x}$ with respect to x_i. Then*

$$\begin{aligned}&\left|f(\mathbf{x} + b\mathbf{e}_i) - f(x) - b\nabla_i f(\mathbf{x}) - \dots - \frac{b^j}{j!}\nabla_i^j f(\mathbf{x})\right|\\ &\le \frac{M_j}{(j+1)!}|b|^{j+1}, \forall \mathbf{x} \in \mathbb{R}^d, \text{ and } b \in \mathbb{R}.\end{aligned} \tag{9}$$

Let a be the sampling radius. The following theorem states how to estimate the coordinate gradient $\nabla_i f(\cdot)$ of a M_j coordinate j-smooth function f by $p = \max\{2(j-1), 2\}$ queries at $\mathbf{x} + \frac{p}{2}a\mathbf{e}_i, \dots, \mathbf{x} + a\mathbf{e}_i, \mathbf{x} - a\mathbf{e}_i, \dots, \mathbf{x} - \frac{p}{2}a\mathbf{e}_i$, and provides the error bound.

Theorem 1 (multi-point coordinate gradient estimator) *Assume f is M_j coordinate j-smooth for some $j \in \mathbb{Z}^+$. Let $p = \max\{2(j-1), 2\}$ and $m = \frac{p}{2}$. Define the p-point coordinate gradient estimator of f with respect to some $i \in [d]$ as*

$$\begin{aligned}\tilde{\nabla}_i f(\mathbf{x}) := {}& C_{\frac{p}{2}} f(\mathbf{x} + \frac{p}{2}ae_i) + \dots + C_1 f(\mathbf{x} + ae_i)\\ &- C_1 f(\mathbf{x} - ae_i) - \dots - C_{\frac{p}{2}} f(\mathbf{x} - \frac{p}{2}ae_i),\end{aligned} \tag{10}$$

where

$$\begin{bmatrix} C_1 \\ C_2 \\ \vdots \\ C_{\frac{p}{2}} \end{bmatrix} = \begin{bmatrix} 1 & 2 & \cdots & \frac{p}{2} \\ 1 & 2^3 & \cdots & (\frac{p}{2})^3 \\ & & \vdots & \\ 1 & 2^{p-1} & \cdots & (\frac{p}{2})^{p-1} \end{bmatrix}^{-1} \begin{bmatrix} \frac{1}{2a} \\ 0 \\ \vdots \\ 0 \end{bmatrix}.$$

Then we have the following error bound

$$|\tilde{\nabla}_i f(\mathbf{x}) - \nabla_i f(\mathbf{x})| \le \sum_{q=1}^{m} |C_q| \frac{M_j q^{j+1}}{(j+1)!} a^{j+1}. \tag{11}$$

Remark 4 *Theorem 1 implies that if f is M_1 coordinate 1-smooth (which holds if f is M_1-smooth in the standard notion) or M_2 coordinate 2-smooth, then the coordinate gradient estimator given in* (10) *corresponds to the central difference $\tilde{\nabla}_i f(\mathbf{x}) = \frac{1}{2a}(f(\mathbf{x} + ae_i) - f(\mathbf{x} - ae_i))$, with error bounds of $\frac{M_1}{2}a$ and $\frac{M_2}{6}a^2$ respectively, because $C_1 = \frac{1}{2a}$.*

Algorithm 1: Zeroth-order inexact augmented Lagrangian method (ZO-iALM) for (1)

Initialization: choose $\mathbf{x}^0 \in \mathrm{dom}(f_0), \mathbf{y}^0 = \mathbf{0}$, $\beta_0 > 0$ and $\sigma > 1$
for $k = 0, 1, \ldots,$ **do**
  Let $\beta_k = \beta_0 \sigma^k$, $\phi(\cdot) = \mathcal{L}_{\beta_k}(\cdot, y^k) - h(\cdot)$, and

$$\begin{aligned} \hat{\rho}_k &= \rho_0 + \bar{L}\|\mathbf{y}^k\| + \beta_k \rho_c, \\ \hat{L}_k &= L_0 + \bar{L}\|\mathbf{y}^k\| + \beta_k L_c. \end{aligned} \tag{13}$$

  $\mathbf{x}^{k+1} \leftarrow$ ZO-iPPM$(\phi, h, \mathbf{x}^k, \hat{\rho}_k, \hat{L}_k, \varepsilon)$
  Update $\mathbf{y}$ by

$$\mathbf{y}^{k+1} = \mathbf{y}^k + w_k \mathbf{c}(\mathbf{x}^{k+1}). \tag{14}$$

subroutine ZO-iPPM$(\phi, \psi, \mathbf{x}^0, \rho, L_\phi, \varepsilon)$
  for $t = 0, 1, \ldots,$ **do**
    Let $G(\cdot) = \phi(\cdot) + \rho\| \cdot - \mathbf{x}^t\|^2$
    Obtain $\mathbf{x}^{t+1}$ by a ZOM such that $\mathrm{dist}(\mathbf{0}, \partial(G + \psi)(\mathbf{x}^{t+1})) \le \frac{\varepsilon}{4}$
    if $2\rho\|\mathbf{x}^{t+1} - \mathbf{x}^t\| \le \frac{\varepsilon}{2}$, **then** return $\mathbf{x}^{t+1}$.

In general, we establish that under the j-th order smoothness assumption for some $j \in \mathbb{Z}^+$, the error of the $\max\{2, 2(j-1)\}$-point coordinate gradient estimator is upper bounded by $O(a^j)$, where a is the sampling radius.

A Novel AL-based ZOM

In this section, we present a novel ZOM for solving (1) under the ALM framework, with each ALM subproblem approximately solved by an inexact proximal point method (iPPM).

The pseudocode of our AL-based ZOM for (1) is shown in Algorithm 1 that uses the following notations

$$B_0 = \max_{\mathbf{x} \in \mathrm{dom}(h)} \max\left\{|f_0(\mathbf{x})|, \|\nabla g(\mathbf{x})\|\right\},$$
$$B_c = \max_{\mathbf{x} \in \mathrm{dom}(h)} \|J_c(\mathbf{x})\|, \tag{12a}$$
$$B_i = \max_{\mathbf{x} \in \mathrm{dom}(h)} \max\left\{|c_i(\mathbf{x})|, \|\nabla c_i(\mathbf{x})\|\right\}, \forall\, i \in [l], \tag{12b}$$
$$\bar{B}_c = \sqrt{\textstyle\sum_{i=1}^l B_i^2}, \quad \bar{L} = \sqrt{\textstyle\sum_{i=1}^l L_i^2},$$
$$\rho_c = \sum_{i=1}^l B_i \rho_i, \quad L_c = \sum_{i=1}^l B_i L_i + B_i^2, \tag{12c}$$

where $\{\rho_i\}$ and $\{L_i\}$ are given in Assumption 1.

Notice that Algorithm 1 follows the standard framework of the ALM and uses ZO-iPPM to solve each ALM subproblem. In principle, one can use any ZOM as a subroutine to solve ALM subproblems, such as ZO-AdaMM (Chen et al. 2019) and ZO-proxSGD (Ghadimi, Lan, and Zhang 2016). However, the use of ZO-iPPM (together with our developed zeroth-order accelerated proximal coordinate update) not only leads to best known complexity results, but also gives better numerical performance, as we show in Section 11.

The proposed ZO-iALM is triple-looped. An algorithm with fewer loops would be preferable. However, we are not aware of any existing simpler ZOMs with the same theoretical guarantees as our method for solving functional constrained black-box optimization. An important future direction is to reduce the number of loops and achieve the same theoretical guarantees. Nevertheless, as we demonstrate in Section 11, our algorithm can be efficiently implemented without much difficulty. Specifically, to have a good practical performance, all parameters except the smoothness constant do not require much tuning at all, and most can be constant across different problems. Even with triple loops, the proposed ZO-iALM performs well numerically. Furthermore, some existing FOMs are also triple-looped and can perform better than double-looped FOMs; see (Li et al. 2021) for example.

The kernel problems that we solve within the iPPM are strongly-convex composite problems. Below, we design a zeroth-order accelerated proximal coordinate update (ZO-APCU) method.

Core subsolver: ZO-APCU

In this subsection, we give our core ZO subsolver, called ZO-APCU, to obtain $\mathbf{x}^{t+1}$ in the ZO-iPPM subroutine. Though ZO-APCU will be used for solving subproblems of our proposed ZOM for (1), it has its own merit and appears to be the first proximal coordinate update method *with acceleration* by only using function values of the smooth part. It solves strongly-convex composite problems in the form of

$$\min_{\mathbf{x} \in \mathbb{R}^d} F(\mathbf{x}) := G(\mathbf{x}) + H(\mathbf{x}), \tag{15}$$

where G is a *black-box* μ-strongly convex and L-smooth function, and H is a *white-box* closed convex function.

We make the following assumptions on G and H.

Assumption 3 (coordinate smooth) *G is M_j coordinate j-smooth, for some $j \in \mathbb{Z}^+$.*

Note that if G is L-smooth, Assumption 3 must hold for $j = 1$ and $M_1 = L$.

Assumption 4 *The function H is coordinate-separable, i.e., $H(\mathbf{x}) = \sum_{i=1}^d H_i(x_i)$, where each $H_i(\cdot)$ is convex.*

The pseudocode of ZO-APCU is shown in Algorithm 2, with its equivalent and efficient implementation (which avoids full-dimensional vector operations) given in the Appendix. The design is inspired from the APCG method in (Lin, Lu, and Xiao 2014). A zeroth-order accelerated random search (ZO-ARS) method has been designed in (Nesterov and Spokoiny 2017) to solve (15). Although our ZO-APCU has the same-order query complexity as ZO-ARS, it significantly outperforms ZO-ARS in practice, because ZO-APCU exploits the coordinate-structure and uses more accurate coordinate gradient estimator.

In Algorithm 2, to obtain the required (coordinate) gradient estimates, we use the p-point coordinate gradient estimator defined in (10), where $p = \max\{2(j-1), 2\}$. Let

$$E_i = \sum_{q=1}^m |C_q| \frac{M_j q^{j+1}}{(j+1)!} a^{j+1}, \forall i \in [d];\ E = \sqrt{\sum_{i=1}^d E_i^2}, \tag{16}$$

Algorithm 2: Zeroth-order accelerated proximal coordinate update for (15): ZO-APCU$(G, H, \mu, L, \varepsilon)$

Input: $\mathbf{x}^0 \in \mathrm{dom}(H)$, tolerance ε, smoothness L, strong convexity μ, and epoch length l.
Initialization: $\mathbf{z}^0 = \mathbf{x}^0, \alpha = \frac{1}{d}\sqrt{\frac{\mu}{L}}$
for $k = 0, 1, \ldots$ **do**
  Let $y^k = \frac{\mathbf{x}^k + \alpha \mathbf{z}^k}{1+\alpha}$
  Sample $i_k \in [d]$ uniformly; compute $\tilde{\nabla}_{i_k} G(\mathbf{y}^k)$ such that $\|\tilde{\nabla}_{i_k} G(\mathbf{y}^k) - \nabla_{i_k} G(\mathbf{y}^k)\| \leq E_{i_k}$.
  Compute $\mathbf{z}^{k+1} = \arg\min_{\mathbf{x}\in\mathbb{R}^d} \{ \frac{dL\alpha}{2} \|\mathbf{x} - (1-\alpha)\mathbf{z}^k - \alpha \mathbf{y}^k\|^2 + \langle \tilde{\nabla}_{i_k} G(\mathbf{y}^k), \mathbf{x}_{i_k} - \mathbf{y}^k_{i_k} \rangle + H_{i_k}(\mathbf{x}_{i_k}) \}$.
  $\mathbf{x}^{k+1} = \mathbf{y}^k + d\alpha(\mathbf{z}^{k+1} - \mathbf{z}^k) + d\alpha^2(\mathbf{z}^k - \mathbf{y}^k)$.
  if $k+1 \equiv 0 \pmod{l}$ **then**
    Compute $\tilde{\nabla} G(\mathbf{x}^{k+1})$ such that $\|\tilde{\nabla} G(\mathbf{x}^{k+1}) - \nabla G(\mathbf{x}^{k+1})\| \leq E$
    $\hat{\mathbf{x}}^{k+1} = \arg\min_{\mathbf{x}\in\mathbb{R}^d} \{ \langle \tilde{\nabla} G(\mathbf{x}^{k+1}), \mathbf{x} - \mathbf{x}^{k+1} \rangle + \frac{L}{2}\|\mathbf{x} - \mathbf{x}^{k+1}\|^2 + H(\mathbf{x}) \}$
    Return $\hat{\mathbf{x}}^{k+1}$ and **stop** if $\mathrm{dist}\big(\mathbf{0}, \tilde{\nabla} G(\hat{\mathbf{x}}^{k+1}) + \partial H(\hat{\mathbf{x}}^{k+1})\big) \leq \frac{3\varepsilon}{4}$.

where $m = \frac{p}{2}$ and a is the sampling radius. By Theorem 1, E and E_i are upper bounds of the gradient estimation errors for $\nabla G(\cdot)$ and $\nabla_i G(\cdot)$. Let $\bar{\varepsilon} = \frac{\mu}{512L}\varepsilon^2$. We choose $a > 0$ and p such that the error bounds E and $\{E_i\}_{i=1}^d$ in (16) satisfy

$$2L\sqrt{\frac{2ED}{\mu}} + E \leq \frac{\varepsilon}{4}, \quad ED + \sum_{i=1}^{d} E_i D_i \leq \frac{\bar{\varepsilon}}{2}. \qquad (17)$$

Complexity Results

In this subsection, we establish the total query complexity result of Algorithm 1. We first show that the core subsolver ZO-APCU can produce $\mathbf{x}^{t+1}$ desired in the ZO-iPPM subroutine. The theorem below gives the complexity result of ZO-APCU to produce an expected ε-stationary point of (15). The proof is highly nontrivial and given in the appendix.

Theorem 2 *Let $\{\mathbf{x}^k\}, \{\hat{\mathbf{x}}^k\}$ be generated from Algorithm 2. Suppose the gradient error bounds E and $\{E_i\}_{i=1}^d$ satisfy* (17). *Let $\bar{\varepsilon} = \frac{\mu}{512L^2}\varepsilon^2$. Then $T = \left\lceil d\sqrt{\frac{L}{\mu}} \log \frac{2(F(\mathbf{x}^0) - F^*) + \mu\|\mathbf{x}^0 - \mathbf{x}^*\|^2}{\bar{\varepsilon}} \right\rceil$ iterations of ZO-APCU suffice to generate $\hat{\mathbf{x}}^T$ satisfying $\mathbb{E}[\mathrm{dist}(\mathbf{0}, \partial F(\hat{\mathbf{x}}^T))] \leq \varepsilon$.*

Theorem 2 only guarantees that the output $\hat{\mathbf{x}}^T$ nearly satisfies the stationarity condition *in expectation*. In order to show the complexity results of Algorithm 1, we need, in Line 4 of ZO-iALM, the iterate $\mathbf{x}^{k+1}$ obtained from ZO-iPPM *deterministically* satisfies the near-stationarity condition of $\mathcal{L}_{\beta_k}(\cdot, y^k)$ so that we can uniformly bound the AL function at the generated iterates. For this technical reason, we will require the output from Algorithm 2 to satisfy the near-stationarity condition *deterministically* instead of in an expectation sense. Theorem 3 below serves as a bridge to convert deterministic iteration bound until expected convergence to expected iteration bound until deterministic convergence, by only sacrificing a log factor in the iteration bound. The result is not difficult to prove but is essential in our complexity analysis of ZO-iALM.

Theorem 3 (expected complexity) *For a sequence of nonnegative random numbers $\{q_k\}_{k=1}^{\infty}$, suppose $\mathbb{E}[q_k] \leq C\eta^k, \forall k \geq 1$ for some $\eta \in (0, 1)$ and $C > 0$. Given $\varepsilon > 0$, define $K(\varepsilon) = \min_{k\in\mathbb{Z}^+}\{k : q_k \leq \varepsilon\}$. Then $\mathbb{E}[K(\varepsilon)] \leq \frac{2-\eta}{1-\eta} \log \frac{C}{\varepsilon(1-\eta)^2} + 3 - \eta$.*

By Theorem 2 (and its proof for linear convergence) and Theorem 3, we have the following expected iteration complexity result of Algorithm 2 until deterministic convergence.

Corollary 1 *Under the same assumptions as Theorem 2, T iterations of ZO-APCU are enough to generate $\hat{\mathbf{x}}^T$ satisfying $\mathrm{dist}(\mathbf{0}, \partial F(\hat{\mathbf{x}}^T)) \leq \varepsilon$, where $\mathbb{E}[T] = \tilde{O}\left(d\sqrt{\frac{L}{\mu}}\right)$.*

Relying on Corollary 1, the next theorem gives the complexity result of the subroutine ZO-iPPM applied on the nonconvex composite problem

$$\Phi^* = \min_{\mathbf{x}\in\mathbb{R}^d} \big\{ \Phi(\mathbf{x}) := \phi(\mathbf{x}) + \psi(\mathbf{x}) \big\}, \qquad (18)$$

where ϕ is a *black-box* L_ϕ-smooth and ρ-weakly convex function, and ψ is a *white-box* closed convex function.

Theorem 4 *Suppose Φ^* in* (18) *is finite. Then the subroutine ZO-iPPM in Algorithm 1 must stop within T iterations, where $T = \left\lceil \frac{32\rho}{\varepsilon^2}(\Phi(\mathbf{x}^0) - \Phi^*) \right\rceil$. The output $\mathbf{x}^T$ must satisfy $\mathrm{dist}(\mathbf{0}, \partial\Phi(\mathbf{x}^T)) \leq \varepsilon$. In addition, if $\mathrm{dom}(\psi)$ has diameter $D_\psi < \infty$ and ZO-APCU is applied to find each $\mathbf{x}^{t+1}$ in ZO-iPPM, then the expected total query complexity is $\tilde{O}\left(\frac{d\sqrt{\rho L_\phi}}{\varepsilon^2} [\Phi(\mathbf{x}^0) - \Phi^*] \log \frac{D_\psi}{\varepsilon} \right)$.*

Now we are ready to establish the query complexity of the proposed ZO-iALM. Due to the difficulty of the possibly nonconvex constraints, a certain regularity condition must be made in order to guarantee (near) feasibility in a polynomial time. Following (Li et al. 2021; Lin, Ma, and Xu 2019; Sahin et al. 2019) that study FOMs, we assume the following regularity condition on (1).

Assumption 5 (regularity) *There is some $v > 0$ such that for any $k \geq 1$,*

$$v\|\mathbf{c}(\mathbf{x}^k)\| \leq \mathrm{dist}\left(-J_c(\mathbf{x}^k)^\top \mathbf{c}(\mathbf{x}^k), \frac{\partial h(\mathbf{x}^k)}{\beta_{k-1}} \right). \qquad (19)$$

Remark 5 *Notice that we only need the existence of v in Assumption 5 but do not need to know its value in our algorithm. The assumption ensures that a near-stationary point of the AL function is near feasible. In (Li et al. 2021), the regularity condition is proven to hold for all affine-equality constrained problems possibly with either an additional polyhedral or ball constraint set. Moreover, several nonconvex examples satisfying Assumption 5 are given in (Lin, Ma, and Xu 2019; Sahin et al. 2019).*

With Assumption 5, we can simply solve a quadratic-penalty problem of (1) *with a large enough penalty parameter,* *in order to find a near-KKT point of* (1)*. However, this approach is numerically much slower than the iALM framework in Algorithm 1; see the tests in (Li et al. 2021) for example.*

Remark 6 *To solve the nonconvex constrained problem* (1)*, a few existing works about FOMs have made key assumptions different from Assumption 5. For example, the uniform Slater's condition was assumed in (Ma, Lin, and Yang 2020), and a strong MFCQ condition was assumed in (Boob, Deng, and Lan 2019). These assumptions are neither strictly stronger nor strictly weaker than Assumption 5.*

The theorem below gives the total query complexity of ZO-iALM with general dual step sizes.

Theorem 5 (total complexity of ZO-iALM) *Suppose that Assumptions 1, 2, and 5 hold. In Algorithm 1, for some fixed* $q \in \mathbb{Z}^+ \cup \{0\}$ *and* $M > 0$*, let* $w_k = \frac{M(k+1)^q}{\|\mathbf{c}(\mathbf{x}^{k+1})\|}, \forall k \geq 0$*. Then given* $\varepsilon > 0$*, Algorithm 1 can produce an* ε*-KKT solution of* (1) *with* $\tilde{O}(d\varepsilon^{-3})$ *queries to* g *and* $\mathbf{c}$ *in expectation, by using Algorithm 2 to find each* $\mathbf{x}^{t+1}$ *in ZO-iPPM. In addition, if* $\mathbf{c}(\mathbf{x}) = \mathbf{A}\mathbf{x} - \mathbf{b}$*, then* $\tilde{O}(d\varepsilon^{-\frac{5}{2}})$ *queries in expectation are enough for Algorithm 1 to produce an* ε*-KKT solution of* (1)*.*

Remark 7 *The results in Theorem 5 are novel. To the best of our knowledge, they are the first such results for ZOMs on solving functional constrained black-box optimization. The order-dependence on* ε *matches with the best-known results for FOMs on solving nonconvex composite optimization with convex or nonconvex constraints, e.g., see (Lin, Ma, and Xu 2019; Melo, Monteiro, and Wang 2020b; Li et al. 2021). For the affine-constrained case, we conjecture that the* $\tilde{O}(d\varepsilon^{-\frac{5}{2}})$ *query complexity may be reduced to* $\tilde{O}(d\varepsilon^{-2})$ *if the nonsmooth term* h *has some special structure.*

Numerical Results

In this section, we conduct numerical experiments to demonstrate the performance of our proposed ZO-iALM. We consider the problem of resource allocation in sensor networks and the adversarial example generation problem. All the tests were performed in MATLAB 2019b on a Macbook Pro with 4 cores and 16GB memory. Due to the page limitation, we put additional numerical experiments in the Appendix. They are on nonconvex linearly-constrained quadratic programs (LCQP), on the unconstrained strongly-convex quadratic programs (USCQP) to test the core solver ZO-APCU, and on the logistic regression to test different multi-point coordinate gradient estimators. We emphasize here that the proposed ZO-APCU requires significantly fewer queries to reach a near-stationary point to the USCQP problem compared to a few existing methods, and that the use of more points in coordinate gradient estimator can lead to higher accuracy.

Resource Allocation in Sensor Networks

In this subsection, we test our proposed ZO-iALM on the resource allocation problem in sensor networks (Liu et al. 2016). The problem aims at minimizing the estimation error of a random vector with a Gaussian prior probability density function, subject to a constraint on the total number of sensor activations. It can be formulated as

$$\begin{aligned} &\min_{\mathbf{w}\in\mathbb{R}^d} \operatorname{tr}(\Sigma^{-1} + \mathbf{H}^\top(\mathbf{w}\mathbf{w}^\top \circ \mathbf{R}^{-1})\mathbf{H})^{-1}, \\ &\text{s.t. } \mathbf{1}^\top\mathbf{w} \leq s, \mathbf{w} \in \{0,1\}^d, \end{aligned} \tag{20}$$

where each $w_i \in \{0,1\}$ denotes whether the ith sensor is selected, $\mathbf{H} \in \mathbb{R}^{d\times d}$ is the observation matrix, $\Sigma \in \mathbb{R}^{d\times d}$ is the MSE source statistics, and $\mathbf{R} \in \mathbb{R}^{d\times d}$ is the noise covariance matrix. We assume that Σ and $\mathbf{R}$ are symmetric, and $\mathbf{R}$ has small off-diagonal entries. Details of the formulation (20) can be found in (Liu et al. 2016).

ZO optimization methods have been applied in the literature to problem (20), in order to avoid the involved first-order gradient computation (Liu et al. 2018). The use of a ZO solver enables the design of resource management with least prior knowledge, e.g., without having access to the sensing model information encoded in $\mathbf{H}$. The constraint $\mathbf{w} \in \{0,1\}^d$ is combinatorial. Below, we rewrite the 0-1 constraint to $\mathbf{w}^2 - \mathbf{w} = \mathbf{0}$ and also incorporate the constraint $\mathbf{1}^\top\mathbf{w} \leq s$ into the objective by introducing a (fixed) multiplier $\lambda > 0$. More precisely, we apply our ZO-iALM to the problem:

$$\begin{aligned} &\min_{\mathbf{w}\in\mathbb{R}^d} \operatorname{tr}(\Sigma^{-1} + \mathbf{H}^\top(\mathbf{w}\mathbf{w}^\top \circ \mathbf{R}^{-1})\mathbf{H})^{-1} + \lambda\mathbf{1}^\top\mathbf{w}, \\ &\text{s.t. } \mathbf{w}^2 - \mathbf{w} = \mathbf{0}. \end{aligned} \tag{21}$$

Since no existing ZOMs are able to handle nonconvex constrained problems, we compare the proposed ZO-iALM to two other methods that replace our ZO-iPPM subroutine with ZO-AdaMM (Chen et al. 2019) and ZO-ProxSGD (Ghadimi, Lan, and Zhang 2016) respectively.

We set $d = 80$, $\lambda = 0.5$, and $\varepsilon = 0.5$. Following (Liu et al. 2016), we construct $\mathbf{H} = \frac{1}{2}(\bar{\mathbf{H}} + \bar{\mathbf{H}}^\top)$ with each entry of $\bar{\mathbf{H}} \in \mathbb{R}^{d\times d}$ generated from the uniform distribution $\mathcal{U}(0,1)$, $\mathbf{R} = (\frac{1}{2}(\bar{\mathbf{R}} + \bar{\mathbf{R}}^\top))^{-1}$ with each entry of $\bar{\mathbf{R}} \in \mathbb{R}^{d\times d}$ generated from $\mathcal{U}(0, 10^{-3})$, and $\Sigma = \mathbf{I}$. In each call to the ZO-iPPM subroutine, we set the smoothness parameter to $\hat{L}_k = 50 + 0.3\beta_k$. We tune the parameters of ZO-AdaMM to $\alpha = 1, \beta_1 = 0.75, \beta_2 = 1$, and fix the step size to 0.01 in ZO-ProxSGD. For each method, we choose $a = 10^{-6}$ as the sampling radius and $w_k = \frac{1}{\|\mathbf{c}(\mathbf{x}^k)\|}$ as the dual step size.

In Figure 1, we compare the primal residual trajectories of the proposed ZO-iALM, and the iALM with subroutine ZO-AdaMM in (Chen et al. 2019) and ZO-ProxSGD in (Ghadimi, Lan, and Zhang 2016). The dual residuals by all compared methods are below the error tolerance ε at the end of each outer loop. In Table 5 in the supplementary material, we also report the primal residual, dual residual, running time (in seconds), and the query count, shortened as `pres`, `dres`, `time`, and `#Obj`, for each method. From the results, we conclude that the proposed ZO-iALM with any of the three subroutines is able to reach an ε-KKT point to the resource allocation problem (21). Moreover, the proposed ZO-iPPM subroutine requires fewer queries than other compared ZOMs to find a specified-accurate stationary point to the nonconvex subproblems.

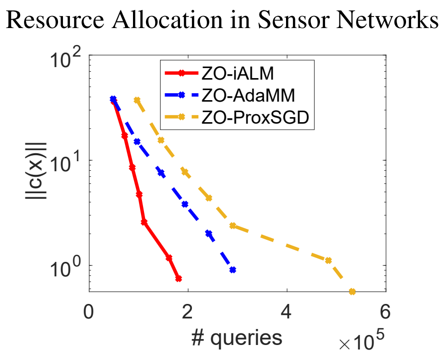

Figure 1: Comparison of iALM on solving (21) with different subroutines: the proposed ZO-iALM, ZO-AdaMM in (Chen et al. 2019), and ZO-ProxSGD in (Ghadimi, Lan, and Zhang 2016). The plots show primal residuals. The markers denote the outer iterations in iALM. Dual residuals for all methods are below the given tolerance ε.

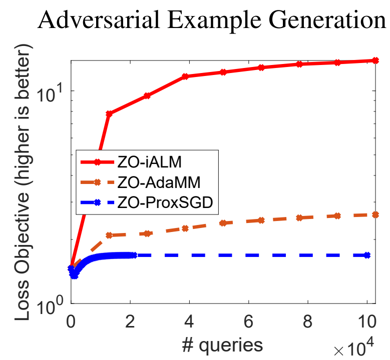

Figure 2: Comparison of iALM on solving (23) with different subroutines: the proposed ZO-iALM, ZO-AdaMM in (Chen et al. 2019), and ZO-ProxSGD in (Ghadimi, Lan, and Zhang 2016). The plots show the loss objective that we attack under the same L_0 and L_∞ constraints.

Adversarial Example Generation

The problem of adversarial example generation for a black-box regression model (Liu et al. 2020b) under both L_0 and L_∞-norm constraints can be formulated as

$$\max_{\|\Delta\|_\infty \le \varepsilon_\infty} f_\theta(\mathbf{x}+\Delta), \text{ s.t. } \|\Delta\|_0 \le \varepsilon_0, \tag{22}$$

where $f_\theta(\cdot)$ is a loss function of a black-box regression model parameterized by θ that is trained over the dataset $\mathbf{x} = [\mathbf{x}_1^\top; \ldots; \mathbf{x}_m^\top] \in \mathbb{R}^{m\times d}$, $\Delta \in \mathbb{R}^d$ is the data perturbation, and $\mathbf{x}+\Delta$ denotes adding Δ to each $\mathbf{x}_i$.

The constraint $\|\Delta\|_0 \le \varepsilon_0$ is combinatorial. To relax it to a continuous one, we introduce a binary vector $\hat{M}$ as a mask and put the constraint onto $\hat{M}$. More precisely, replace Δ in (22) by $\hat{M} \circ \Delta$, where $\circ$ denotes the Hadamard (component-wise) product. Then the constraint $\|\Delta\|_0 \le \varepsilon_0$ is relaxed to $\hat{M}_i \in \{0,1\}, \forall\, i$ and $\mathbf{1}^\top \hat{M} \le \varepsilon_0$. By further incorporating the constraint $\mathbf{1}^\top \hat{M} \le \varepsilon_0$ into the objective by introducing a (fixed) multiplier $-\lambda < 0$ and rewrite $\hat{M}_i \in \{0,1\}, \forall\, i$ into $\hat{M}^2 - \hat{M} = \mathbf{0}$, where $\hat{M}^2$ denotes the component-wise square of $\hat{M}$, we have the following reformulation:

$$\max_{\substack{\hat{M},\Delta\in\mathbb{R}^d\\ \|\Delta\|_\infty\le\varepsilon_\infty}} f_\theta(\mathbf{x}_0+\hat{M}\circ\Delta) - \lambda\mathbf{1}^\top\hat{M}, \text{ s.t. } \hat{M}^2-\hat{M}=\mathbf{0}. \tag{23}$$

We test the proposed ZO-iALM on the adversarial example generation problem (23). In the test, we use the ovarian cancer dataset (Conrads et al. 2004; Petricoin III et al. 2002) that are from $m = 216$ patients. Each data point has $d = 4{,}000$ features and a label indicating whether the corresponding patient has ovarian cancer. We first use MATLAB's built-in lasso function (with $\lambda = 0.01$) to train a LASSO regression model parameterized by θ. With the trained model, we treat the regression loss $f_\theta(\cdot)$ as a ZO oracle and perform black-box attack on it. Let $\mathbf{x} \in \mathbb{R}^{m\times d}$ denote the data matrix. We then solve the ZO formulation (23) to find an adversarial perturbation $M \circ \Delta$ to each row of $\mathbf{x}$ that near-maximally increases the regression loss $f_\theta(\cdot)$. In (23), we set $\lambda = 0.01$ and $\varepsilon_\infty = 0.1$. Due to the large variable dimension, we set $\varepsilon = 1$ in stopping conditions.

The same as the previous test, we compare the proposed ZO-iALM to two other methods that replace our ZO-iPPM subroutine with ZO-AdaMM (Chen et al. 2019) and ZO-ProxSGD (Ghadimi, Lan, and Zhang 2016) respectively. In each method, we set $a = 10^{-6}$ as the sampling radius and $w_k = \frac{1}{\|\mathbf{c}(\mathbf{x}^k)\|}$ as the dual step size.

Let $(\hat{M}, \Delta)$ be one iterate obtained by one method on solving (23). Then $\tilde{\Delta} \leftarrow \hat{M} \circ \Delta$ is the data perturbation. To recover the solution to (22), we project $\tilde{\Delta}$ to the set $\{\Delta : \|\Delta\|_0 \le 20, \|\Delta\|_\infty \le 0.1\}$. In Figure 2, we plot the trajectory of the loss objective f_θ by all methods at the processed iterates of perturbed data. From the results, we see that the data perturbation created by the proposed ZO-iALM increases the loss function faster (namely, creates more successful attacks) than other compared methods.

Conclusion

In this paper, we propose a novel zeroth-order inexact augmented Lagrangian method (ZO-iALM) to solve black-box optimization problems that involve a composite (i.e., smooth+nonsmooth) objective and nonlinear functional constraints. The kernel subproblems that we solve during the ZO-iALM are black-box strongly-convex composite problems with coordinate structure. To most efficiently solve these subproblems, we design a zeroth-order accelerated proximal coordinate update (ZO-APCU) method. In addition, in order to be able to produce high-accurate solutions, we give a new multi-point coordinate gradient estimator and use it in our designed ZO-APCU. All our proposed zeroth-order methods achieve similar-order complexity results as the best-known results obtained by first-order methods, with a difference up to a factor of variable dimension. Besides the novel and best theoretical results, our proposed ZO-iALM can also perform well numerically, which is demonstrated by experiments on practical machine learning tasks and classical optimization problems.

Acknowledgements

This work was supported by the Rensselaer-IBM AI Research Collaboration (http://airc.rpi.edu), part of the IBM AI Horizons Network (http://ibm.biz/AIHorizons). The work of Y. Xu is partly supported by NSF Award 2053493 and the RPI-IBM AIRC faculty fund.

References

Berahas, A. S.; Cao, L.; Choromanski, K.; and Scheinberg, K. 2021. A theoretical and empirical comparison of gradient approximations in derivative-free optimization. *Foundations of Computational Mathematics*, 1–54.

Boob, D.; Deng, Q.; and Lan, G. 2019. Stochastic first-order methods for convex and nonconvex functional constrained optimization. *arXiv preprint arXiv:1908.02734*.

Chen, P.-Y.; Zhang, H.; Sharma, Y.; Yi, J.; and Hsieh, C.-J. 2017. ZOO: Zeroth Order Optimization Based Black-box Attacks to Deep Neural Networks Without Training Substitute Models. In *ACM Workshop on Artificial Intelligence and Security*, 15–26.

Chen, X.; Liu, S.; Xu, K.; Li, X.; Lin, X.; Hong, M.; and Cox, D. 2019. Zo-adamm: Zeroth-order adaptive momentum method for black-box optimization. *arXiv preprint arXiv:1910.06513*.

Conrads, T. P.; Fusaro, V. A.; Ross, S.; Johann, D.; Rajapakse, V.; Hitt, B. A.; Steinberg, S. M.; Kohn, E. C.; Fishman, D. A.; Whitely, G.; et al. 2004. High-resolution serum proteomic features for ovarian cancer detection. *Endocrine-related cancer*, 11(2): 163–178.

Duchi, J. C.; Jordan, M. I.; Wainwright, M. J.; and Wibisono, A. 2015. Optimal rates for zero-order convex optimization: The power of two function evaluations. *IEEE Transactions on Information Theory*, 61(5): 2788–2806.

Flaxman, A. D.; Kalai, A. T.; and McMahan, H. B. 2004. Online convex optimization in the bandit setting: gradient descent without a gradient. *arXiv preprint cs/0408007*.

Ghadimi, S.; Lan, G.; and Zhang, H. 2016. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Mathematical Programming*, 155(1-2): 267–305.

Kiefer, J.; Wolfowitz, J.; et al. 1952. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, 23(3): 462–466.

Li, F.; and Qu, Z. 2021. An inexact proximal augmented Lagrangian framework with arbitrary linearly convergent inner solver for composite convex optimization. *Mathematical Programming Computation*, 1–62.

Li, Z.; Chen, P.-Y.; Liu, S.; Lu, S.; and Xu, Y. 2021. Rate-improved inexact augmented Lagrangian method for constrained nonconvex optimization. In *International Conference on Artificial Intelligence and Statistics*, 2170–2178. PMLR.

Li, Z.; and Xu, Y. 2021. Augmented Lagrangian–Based First-Order Methods for Convex-Constrained Programs with Weakly Convex Objective. *INFORMS Journal on Optimization*, 3(4): 373–397.

Lian, X.; Zhang, H.; Hsieh, C.-J.; Huang, Y.; and Liu, J. 2016. A comprehensive linear speedup analysis for asynchronous stochastic parallel optimization from zeroth-order to first-order. *arXiv preprint arXiv:1606.00498*.

Lin, Q.; Lu, Z.; and Xiao, L. 2014. An accelerated proximal coordinate gradient method. In *Advances in Neural Information Processing Systems*, 3059–3067.

Lin, Q.; Ma, R.; and Xu, Y. 2019. Inexact Proximal-Point Penalty Methods for Non-Convex Optimization with Non-Convex Constraints. *arXiv preprint arXiv:1908.11518*.

Liu, S.; Chen, J.; Chen, P.-Y.; and Hero, A. 2018. Zeroth-order online alternating direction method of multipliers: Convergence analysis and applications. In *International Conference on Artificial Intelligence and Statistics*, 288–297. PMLR.

Liu, S.; Chen, P.-Y.; Kailkhura, B.; Zhang, G.; Hero, A.; and Varshney, P. K. 2020a. A Primer on Zeroth-Order Optimization in Signal Processing and Machine Learning. *IEEE Signal Processing Magazine*.

Liu, S.; Chepuri, S. P.; Fardad, M.; Maşazade, E.; Leus, G.; and Varshney, P. K. 2016. Sensor selection for estimation with correlated measurement noise. *IEEE Transactions on Signal Processing*, 64(13): 3509–3522.

Liu, S.; Lu, S.; Chen, X.; Feng, Y.; Xu, K.; Al-Dujaili, A.; Hong, M.; and O'Reilly, U.-M. 2020b. Min-max optimization without gradients: Convergence and applications to black-box evasion and poisoning attacks. In *International Conference on Machine Learning*, 6282–6293. PMLR.

Ma, R.; Lin, Q.; and Yang, T. 2020. Quadratically Regularized Subgradient Methods for Weakly Convex Optimization with Weakly Convex Constraints. In *International Conference on Machine Learning*, 6554–6564. PMLR.

Melo, J. G.; Monteiro, R. D.; and Wang, H. 2020a. Iteration-complexity of an inexact proximal accelerated augmented Lagrangian method for solving linearly constrained smooth nonconvex composite optimization problems. *Optimization Online*.

Melo, J. G.; Monteiro, R. D.; and Wang, H. 2020b. Iteration-complexity of an inexact proximal accelerated augmented Lagrangian method for solving linearly constrained smooth nonconvex composite optimization problems. *arXiv preprint arXiv:2006.08048*.

Necoara, I.; and Nedelcu, V. 2014. Rate analysis of inexact dual first-order methods application to dual decomposition. *IEEE Transactions on Automatic Control*, 59(5): 1232–1243.

Nedelcu, V.; Necoara, I.; and Tran-Dinh, Q. 2014. Computational complexity of inexact gradient augmented Lagrangian methods: application to constrained MPC. *SIAM Journal on Control and Optimization*, 52(5): 3109–3134.

Nesterov, Y.; and Spokoiny, V. 2017. Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 17(2): 527–566.

Ouyang, Y.; Chen, Y.; Lan, G.; and Pasiliao Jr, E. 2015. An accelerated linearized alternating direction method of multipliers. *SIAM Journal on Imaging Sciences*, 8(1): 644–681.

Petricoin III, E. F.; Ardekani, A. M.; Hitt, B. A.; Levine, P. J.; Fusaro, V. A.; Steinberg, S. M.; Mills, G. B.; Simone, C.; Fishman, D. A.; Kohn, E. C.; et al. 2002. Use of proteomic patterns in serum to identify ovarian cancer. *The lancet*, 359(9306): 572–577.

Sahin, M. F.; Alacaoglu, A.; Latorre, F.; Cevher, V.; et al. 2019. An inexact augmented lagrangian framework for non-convex optimization with nonlinear constraints. In *Advances in Neural Information Processing Systems*, 13943–13955.

Xu, Y. 2021a. First-order methods for constrained convex programming based on linearized augmented Lagrangian function. *INFORMS Journal on Optimization*, 3(1): 89–117.

Xu, Y. 2021b. Iteration complexity of inexact augmented lagrangian methods for constrained convex programming. *Mathematical Programming*, 185(1): 199–244.